

Approximate Revenue Maximization in Interdependent Value Settings *

Shuchi Chawla
University of Wisconsin - Madison
Computer Sciences Department
shuchi@cs.wisc.edu

Hu Fu
Microsoft Research
New England
hufu@microsoft.com

Anna R. Karlin
University of Washington
Department of Computer Science and Engineering
karlin@cs.washington.edu

August 20, 2014

Abstract

We study revenue maximization in settings where agents' values are interdependent: each agent receives a signal drawn from a correlated distribution and agents' values are functions of all of the signals. We introduce a variant of the generalized VCG auction with reserve prices and random admission, and show that this auction gives a constant approximation to the optimal expected revenue in matroid environments. Our results do not require any assumptions on the signal distributions, however, they require the value functions to satisfy a standard single-crossing property and a concavity-type condition.

*Part of this work was done while the first and third authors were visiting Microsoft Research. The first author is supported in part by NSF grants CCF-1101429 and CCF-1320854, and the third author by NSF grant CCF-1016509.

1 Introduction

Revenue maximization, a.k.a. *optimal auction* design, is a fundamental goal in mechanism design. The most common approach to the problem is Bayesian wherein it is assumed that players' values are drawn from prior distributions known to the designer. The seminal work of Myerson [24] on this topic assumes the agents' values are drawn from independent distributions. In this paper, we consider settings in which this independence assumption is unreasonable. If the value of the item derives from its resale value, such as a house, the values of the agents are likely to be correlated. Yet another example is fashion, whether haute couture or fancy cars, where demand is highly correlated. Likewise, when the value of the item derives from the potential for making money from it, the values are correlated, or perhaps even common. The classical example is an auction for the right to drill for oil in a certain location [36]. Note though that even in settings such as this where the value is common, bidders might have different information about what that value actually is. For example, the value of an oil lease depends on how much oil there actually is, and the different bidders may have access to different assessments about this. Consequently, a bidder might change her own estimate of the value of the oil lease given access to the information another bidder has.

The following model due to Milgrom and Weber [23] has become standard for auction design in these correlated and common value settings, generally known as *interdependent settings*: [See also 17, 22]

- Each agent has a real-valued, private *signal*, say s_i for agent i . The set of signals $\mathbf{s} = (s_1, s_2, \dots, s_n)$ is assumed to be drawn from a known, correlated distribution.

The signals are meant to capture the information available to the agents about the item. For example, in the setting of a painting, the signal could be information about the provenance of the painting or private information from experts. In the setting of oil drilling rights, the signals could be information that engineers have about the site based on geologic surveys, etc.

- The *value of the item* to agent i is a function $v_i(\mathbf{s})$ of the information or signals of *all* agents.

For example, in the *mineral rights model* [36], each signal $S_i = V + X_i$ might be a noisy version of the true value V of the oil field (say X_i is $N(0, 1)$), and $v_i(\mathbf{s}) = E(V | s_1, \dots, s_n)$.

Another important case is when $v_i(\mathbf{s}) = s_i + \beta \sum_{j \neq i} s_j$, for some $\beta \leq 1$. This type of valuation function captures settings where an agent's value depends on his own personal appreciation for the item (s_i) and on resale value (as reflected by the total appreciation of the other agents) [24, 16].

Interdependent settings have received attention in the economics community for quite some time [see 17].¹ Perhaps the most notable results here are due to Crémer and McLean [8, 9] [see also 20, 21, 27]. These papers consider single item auctions and show that if the prior distribution from which \mathbf{s} is drawn satisfies a type of full rank condition, and the valuation functions are known and satisfy a single-crossing condition (see Section 2), then the auctioneer can construct an auction that essentially extracts full surplus (that is, has expected revenue equal to $\mathbf{E}[\max_i v_i(S_1, \dots, S_n)]$).

From a practical perspective, the Crémer-McLean auction is problematic in two significant ways: First, the auction is not *ex post* individually rational. That is, after the fact, agents may regret

¹Within the theoretical computer science community, interdependent settings have received far less attention and most of that only in the last few years. [See, e.g. 28, 26, 13, 4, 1, 33, 29, 18].

having participated, and indeed may end up being charged a very large amount to participate. Second, the dependence on the signal distribution is complex in form and can be numerically unstable. Moreover, agents need to know the prior distribution from which the signals are drawn. For these reasons, this result has been criticized as being impractical [e.g., in 22].

These shortcomings have driven much recent progress in the understanding of optimal auctions. Addressing the first issue, i.e., strengthening the requirement for individual rationality so that bidders are guaranteed non-negative utilities under any circumstance, a series of recent papers [32, 5, 6, 35, 18, 19] studies *ex post equilibria*,² culminating in Roughgarden and Talgam-Cohen [29]’s characterization of the expected revenue as the “conditional virtual surplus”. This characterization was used to derive optimal auctions for symmetric matroid settings with various assumptions on the distribution and the value functions (see Section 2 for details). In particular, for the general interdependent setting, the value distribution is required to satisfy a *monotone hazard rate (MHR)* condition, a property that holds for certain unimodal distributions with light tails.

In this work, we design auctions for a broad class of interdependent value settings that obtain approximately optimal revenue in ex post equilibrium in the absence of any distributional assumptions.

Our auctions are variants of the VCG auction with reserve prices. The VCG auction with reserve prices is known to be approximately optimal in independent private value settings under certain assumptions. Myerson’s characterization of the optimal auction implies that VCG with monopoly reserve prices is optimal in the i.i.d. values setting under a matroid feasibility constraint. In non-i.i.d. settings, the optimal auction can be complex and may be impractical. Hartline and Roughgarden [15] initiated the study of VCG auctions with reserve prices, as a model for simple and yet approximately optimal auctions.

For interdependent settings that satisfy a single-crossing condition (a condition also assumed by Roughgarden and Talgam-Cohen [29] and virtually all literature on these settings), Li [18] showed that a generalized notion of the VCG auction with appropriate reserve prices gives an ϵ -approximation to the optimal revenue for matroid settings with generalized MHR distributions. Unfortunately, for more general distributions, such reserve-based VCG auctions perform poorly (see example in Section 1.2 below).

We introduce a new variant of VCG auctions: a random admission phase followed by a VCG auction with reserve prices. We show that this randomization preserves the incentive compatibility and the approximately optimal performance of VCG auctions even when all restrictions on the value distribution are lifted, although we will need a certain condition on the value function. As a special case of our main result, we have:

Theorem 1.1. *The randomized GVCG- L^* auction described in Section 4 is ex post incentive compatible, individually rational, and obtains a constant fraction of the optimal revenue under a matroid feasibility constraint, assuming that the valuation functions are additively separable functions of signals, and satisfy the single-crossing condition.*

Theorem 1.1 in fact holds for a larger class of valuation functions that includes, e.g., concave functions of additively separable functions of signals (see Assumption (A.3) in Section 2). Although Theorem 1.1 can be seen as a simple vs. optimal type of result, we emphasize that, for non-MHR

²Bidding truthfully is an *ex post equilibrium* if an agent does not regret having bid truthfully when the mechanism terminates, given that other agents bid truthfully. In other words, bidding truthfully is a Nash equilibrium at the end. See Section 2.

value distributions, no optimal auctions or approximately optimal ones were previously known, even for selling a single item.

We now describe our results in more detail.

1.1 Correlated Private Values

We first consider the special case where the values ($v_i := v_i(\mathbf{s}) = s_i$) are privately known and drawn from a correlated distribution F . The computational complexity of designing the revenue maximizing mechanism in this setting at ex post Nash equilibrium³ under the assumption of ex post individual rationality has been largely resolved. Papadimitiou and Pierrakos [26] proved that the deterministic revenue maximization is computationally hard even with 3 bidders. In some settings positive results were obtained: Dobzinski et al. [13] showed how to derive truthful-in-expectation mechanisms via linear programming when the joint distribution of values is given explicitly; Roughgarden and Talgam-Cohen [29] give optimal auctions for matroid settings under the assumptions of affiliation and a strong version of regularity.

Lookahead auctions for a single item. Ronen [28] took the approximation approach, and proposed the following intuitive *lookahead* auction: choose the highest bidder as the tentative winner and then run an optimal auction for the highest bidder conditioning on all other bidders' valuations and the fact that this tentative winner's valuation is above all others'. It is not hard to see that this auction is dominant strategy incentive-compatible for correlated bidders. Moreover, with a simple and elegant proof, Ronen showed that this auction gives at least half of the optimal revenue of any single-item auction. Dobzinski et al. [13] later extended these to a more general class of lookahead auctions with better approximation ratios.

Our first result generalizes the lookahead auction to the matroid setting for correlated bidders in the most natural way, and shows that it is still a 2-approximation.

Theorem 1.2. *The lookahead auction in Section 3 gives at least half of the optimal revenue in any matroid setting for bidders with arbitrarily correlated valuations.*

It is interesting to note that the lookahead auction (in both the single-item case and in our generalization for the matroid case) is a special case of *VCG-L auctions*: A VCG-L auction [12, 18] is an auction that first selects as tentative winners the bidders who would be winners in a VCG auction, and then offers a take-it-or-leave-it price to each of these tentative winners, where the price is the higher of the VCG price and a preset reserve price. In the lookahead auction, the reserve price is determined optimally based on the conditional value distribution of the tentative winner, conditioned on others' revealed values and on being the tentative winner; We call this the *conditional monopoly reserve*. In fact, the lookahead auction or VCG-L with conditional monopoly reserves is the optimal among all VCG-L auctions.

1.2 General Interdependent Settings

General interdependent settings are far less understood than private value settings. Indeed, even the more basic problem of maximizing social welfare is not straightforward: It is no longer possible to design dominant-strategy auctions, since an agent's value depends on all the signals, and so if,

³For correlated values, ex post equilibrium and dominant-strategy equilibrium coincide. See Section 2.

say, agent i reports rubbish, agent j might win at a price above his value if he reports truthfully. Hence, the standard approach is to try maximizing welfare at ex post equilibria. Unfortunately, maximizing social welfare in an ex post equilibrium is also provably impossible unless the valuation functions $v_i(\mathbf{s})$ satisfy a so-called *single-crossing condition*. This condition means, in essence, that the influence of the highest bidder's signal on her own value is at least as high as its influence on the other high bidders' values.⁴ When the single-crossing condition holds, there is a generalization of VCG auctions that maximizes efficiency at ex post equilibrium for matroid settings [8, 9, 17, 2, 10, 18]: the mechanism asks agents to reveal their signals, computes values based on the reported signals, and then runs a VCG-type auction on these values.

Returning to the goal of approximate revenue maximization, a generalization of VCG-L with conditional monopoly reserves was proposed by Li [18]. This auction, which we will call *Generalized VCG-L with conditional monopoly reserves* or GVCG-L* for short, is the following:

1. Ask agents to report their signals; compute agents' values.
2. Choose the social welfare-maximizing set W , and determine the prices they would be charged in the generalized VCG auction.
3. Offer each agent i in W a take-it-or-leave-it price which is the maximum of the price from step (2), and the optimal price conditioned on i being in W and conditioned on the signals of all the other agents.

This auction was shown by Li [18] to obtain an expected revenue which approximates the optimal social welfare (and hence optimal revenue as well) for matroid settings under the assumption that the conditional distributions satisfy the monotone hazard rate (MHR) condition. Li [18]'s analysis relies on a property peculiar to MHR distributions, namely that the optimal revenue from a single agent is close to the agent's expected value.

Unfortunately, GVCG-L* does not give any constant approximation to the optimal revenue if the conditional distributions are not MHR, even for single item settings. Consider, for instance, a setting with two agents and one item where the first agent's value is $v_1(s_1, s_2) = s_1$ and the second agent's value is $v_2(s_1, s_2) = s_1 - \epsilon$. In this case, GVCG-L* always serves the first agent, obtaining the same expected revenue as in a single agent auction with this agent alone, because the signal s_2 contains no information and v_2 is completely determined by s_1 (incentive compability prevents the use of v_2 in extracting revenue from the first agent). The optimal auction, on the other hand always serves the second agent obtaining her entire value (because s_1 completely determines v_2). The gap between the revenues of these two auctions can be arbitrarily large if the distribution of s_1 is non-MHR. We note, in particular, that a revenue-maximizing auction must at times sell to an agent whose value is not the largest. Accordingly, we add a stage of random admission at the start of the auction, ensuring that the auction precludes the highest-value agent from winning with constant probability, while still using her signal to help with setting the prices for the other agents. This enables us to overcome the shortcoming of the GVCG-L* auction and to recover revenue from all cases with constant probability.

1.3 Results in Independent Private Value Settings

Ronen [28]'s analysis of the lookahead auction uses first principles without appealing to Myerson's

⁴ This implies that, given signals \mathbf{s}_{-i} , if agent i has the highest value when $s_i = t^*$, then agent i continues to have the highest value for $s_i > t^*$.

virtual values. Taking the same approach allows us to develop new insight into the design of simple approximately-optimal auctions even in *independent private value* settings. In Appendix A we rederive several known simple vs. optimal results [e.g., 15, 3] as quick consequences of the lookahead auction. We single out one step of this in Appendix B, which may be of independent interest. This concerns the VCG-E or *VCG with eager reserves* auction, which differs from the VCG-L auction only in that it removes all bidders bidding below their reserve prices *before* picking the tentative winners. We show that, when the reserve prices are at least the monopoly reserves, VCG-E is always preferable than VCG-L, in terms of both social welfare and revenue (Theorem B.2).

1.4 Other Related Work

We briefly give further details on some of the related work.

The standard auctions. The revenue in Bayes-Nash equilibrium of the standard auctions (English, first price and second price) has been analyzed assuming a strong form of positive correlation known as “affiliation” between the signals [23]. One of the most interesting results here is the “linkage principle” which shows that in the symmetric, affiliated setting, the auctioneer benefits by enabling the maximum amount of information to be revealed. A consequence is that open auctions such as the English auction generally lead to higher prices than closed auctions such as sealed-bid first price auctions.

Characterization of revenue in ex post equilibrium. As mentioned earlier, a series of papers, including [32, 5, 6, 35, 18, 29, 19] [see 29, for a literature review], develop characterization results for ex post equilibrium that show that the expected revenue in ex post equilibrium is equal to the expected “conditional virtual surplus”.

Optimal auctions in ex post equilibrium. The characterization results just mentioned are used to derive Myerson-like optimal auctions in some settings. The most broadly applicable results here are due to Roughgarden and Talgam-Cohen [29] who use their characterization to derive optimal auctions (a) for correlated values in matroid settings, with assumptions of regularity and affiliation, and (b) for the general interdependent case when bidders are symmetric, have affiliated signals, and satisfy a number of other conditions [19]. They also derive the first approximately-optimal prior-independent results under the Lopomo assumptions for matroid feasibility constraints. An auction is *prior independent* if it does not require knowing the distributions of agents’ values [12].

Several of these papers [e.g., 19, 25, 6, 29] focus specifically on the English auction and derive among other things that the English auction with a suitable reserve price is optimal among all ex post IC and ex post IR auctions for symmetric settings with correlated values that satisfy regularity and affiliation.

Bayes-Nash equilibrium versus ex post equilibrium. We refer the reader to Roughgarden and Talgam-Cohen [29] for an excellent discussion of the tradeoffs between BIC mechanism design and ex post IC mechanism design.

2 Preliminaries

Single Dimensional Environments. In a single dimensional auction environment, an auctioneer is offering a service (or goods) to n bidders, but with certain constraints on which subsets of bidders can be simultaneously served. Formally, if we use $[n]$ to denote the set of bidders, then there is a set $\mathcal{I} \subseteq 2^{[n]}$, such that the auctioneer can serve a subset S of bidders simultaneously if and only if S is in \mathcal{I} . For example, in a single item auction, \mathcal{I} consists of all singleton subsets of $[n]$ and the empty set. We say the environment is *downward closed* if $T \in \mathcal{I}$ implies $S \in \mathcal{I}, \forall S \subseteq T$.

A feasibility constraint \mathcal{I} is called a *matroid constraint* if it is downward closed and for all $A, B \in \mathcal{I}$ with $|A| > |B|$, there exists $e \in A$ such that $B \cup \{e\} \in \mathcal{I}$. Matroid settings encompass important types of markets, e.g. digital goods (when all subsets are feasible, i.e., $\mathcal{I} = 2^{[n]}$), k unit auctions (when \mathcal{I} contains all subsets of size at most k , i.e., is the k -uniform matroid), and unit-demand bipartite matching markets (when \mathcal{I} is a transversal matroid).

Signals, valuations and distributions. Each bidder i has a private one-dimensional signal s_i representing the information available to him. The signal s_i is a realization of a random variable S_i . We will assume that the signals $\mathbf{S} = (S_1, \dots, S_n)$ are drawn from a known *correlated* joint distribution denoted by F . We denote the vector $(s_1, \dots, s_{i-1}, s'_i, s_{i+1}, \dots, s_n)$ by (s'_i, \mathbf{s}_{-i}) .

An agent's valuation $v_i(\mathbf{s})$ is a function of all the signals. (In the special case of correlated or independent private values, v_i is equal to the signal s_i and in that case, we do not refer to signals, but rather to the valuations v_i themselves.) These valuation functions are assumed to be common knowledge, and satisfy the following conditions:

- (A.1) For all i , $v_i(\mathbf{s})$ is increasing in each coordinate and strictly increasing in s_i .
- (A.2) The valuations satisfy the *single-crossing condition*: For each $i \neq j$, \mathbf{s}_{-i} , s_i , and s'_i , with $s'_i > s_i$, $v_i(s_i, \mathbf{s}_{-i}) \geq v_j(s_i, \mathbf{s}_{-i})$ implies $v_i(s'_i, \mathbf{s}_{-i}) > v_j(s'_i, \mathbf{s}_{-i})$. In other words, as soon as agent i 's value crosses agent j 's value, then increasing i 's signal continues to keep v_i larger than v_j .
- (A.3) The responsiveness of an agent's value to someone else's signal decreases as the agent's own signal increases. Formally, for all $i \neq j$ and for all \mathbf{s} , $\frac{\partial v_i(\mathbf{s})}{\partial s_j}$ is a non-increasing function of s_i .

Assumptions (A.1) and (A.2) are standard in interdependent value settings. Assumption (A.3) is new. One broad class of value functions for which this assumption holds is *additively separable* valuations, i.e. $v_i(\mathbf{s}) = \sum_j g_{ij}(s_j)$ where each $g_{ij}(\cdot)$ is a non-decreasing functions. More generally, if values are concave functions of additively separable functions of signals, then assumption (A.3) holds.

Auctions and Incentive Compatibility. An *auction* takes as input a signal from each bidder and produces as output a pair of vector functions (x, p) defined on tuples of signals. At the signals (s_1, \dots, s_n) , the *allocation* function, $x_i(s_1, \dots, s_n)$, gives the probability with which bidder i receives the service, and $p_i(s_1, \dots, s_n)$, the *payment* function, denotes the expected payment bidder i makes to the auctioneer. (The randomization here is in the mechanism.)

Bidder i 's *utility* from participating in the auction $(x(\cdot), p(\cdot))$, when the true signals are \mathbf{s} , he reports r_i , and the other bidders report \mathbf{r}_{-i} , is $v_i(\mathbf{s})x_i(\mathbf{r}) - p_i(\mathbf{r})$.

Since bidders' signals are private, they may misreport their signals to gain advantage. In order to incentivize truthtelling, we study auctions that are *incentive compatible*. Three different equilibrium notions, presented in order of decreasing strength, will come up in this paper:

- An auction is called *dominant strategy incentive compatible* (DSIC) if, for all i , true signals (s_1, \dots, s_n) , reported signals (r_1, \dots, r_n) and s'_i :

$$v_i(\mathbf{s})x_i(s_i, \mathbf{r}_{-i}) - p_i(s_i, \mathbf{r}_{-i}) \geq v_i(\mathbf{s})x_i(s'_i, \mathbf{r}_{-i}) - p_i(s'_i, \mathbf{r}_{-i}).$$

- An auction is called *ex post incentive compatible* (ex post IC, for short)⁵ if, for all i , true signals (s_1, \dots, s_n) and s'_i

$$v_i(\mathbf{s})x_i(\mathbf{s}) - p_i(\mathbf{s}) \geq v_i(\mathbf{s})x_i(s'_i, \mathbf{s}_{-i}) - p_i(s'_i, \mathbf{s}_{-i}).$$

- An auction is called *Bayesian incentive compatible* (BIC) if, for all i , s_i and s'_i ,

$$\mathbf{E}[v_i(s_i, \mathbf{S}_{-i})x_i(s_i, \mathbf{S}_{-i}) - p_i(s_i, \mathbf{S}_{-i})] \geq \mathbf{E}[v_i(s_i, \mathbf{S}_{-i})x_i(s'_i, \mathbf{S}_{-i}) - p_i(s'_i, \mathbf{S}_{-i})].$$

In this paper we will focus only on incentive compatible auctions, and so we will use the terms *bids* and *signals*, or, in the correlated case, *valuations*, interchangeably.

We will also focus exclusively on auctions that satisfy *ex post individual rationality*, that is each agent's final utility is always nonnegative.

Social Welfare and Revenue. The *social welfare* for serving a set T of bidders with signals \mathbf{s} is $\sum_{i \in T} v_i(\mathbf{s})$. The expected revenue of an auction is $\mathbf{E}[\sum_i p_i(\mathbf{S})]$.

VCG and Generalized VCG auctions with lazy reserve prices. In a VCG auction [34, 7, 14], the winning set of bidders is the feasible set with the largest sum of valuations. Each of them pays the critical valuation below which he would drop from the winning set. We will be interested in the generalization of VCG that applies in interdependent matroid settings that satisfy the single-crossing condition [8, 9, 17, 2, 10, 18].

Formally, in generalized VCG, the set W of winners is

$$W := \operatorname{argmax}_{T \in \mathcal{I}} \sum_{i \in T} v_i(\mathbf{s}).$$

To define payments, let

$$s_i^* = \inf_{s_i} \{i \in \operatorname{argmax}_{T \in \mathcal{I}} \sum_{j \in T} v_j(s_i, \mathbf{s}_{-i})\}.$$

The *VCG payment* or *VCG threshold* for bidder i is then

$$p_i^{\text{VCG}} := v_i(s_i^*, \mathbf{s}_{-i}).$$

This payment p_i^{VCG} is defined solely by the signals of other bidders, the form of the valuation functions, and the feasibility system \mathcal{I} , and therefore is well defined for all bidders, including those

⁵All ex post IC auctions that we consider will, in fact, satisfy the ex post IC condition even after any internal randomization in the auction.

who lose the auction. For each bidder, the VCG threshold is the value above which he will win the auction, and below which he loses.

By construction, the VCG auction maximizes social welfare. Moreover, for correlated settings, VCG is dominant strategy incentive compatible, and for interdependent settings, generalized VCG is ex post incentive compatible. (See Lemma 2.1.)

The main object of study in this paper is a generalized VCG-type auction with individual reserve prices.

Definition 1 (Generalized VCG Auctions with Lazy Reserve Prices (GVCG-L)). In the generalized VCG auction with *lazy reserve prices* (r_1, \dots, r_n) , agents are asked to report their signals. Given these signals, the mechanism first chooses a tentative subset W of winners as in the GVCG auction, i.e., $W = \operatorname{argmax}_{T \in \mathcal{I}} \sum_{i \in T} v_i(\mathbf{s})$. Then each bidder $i \in W$ is presented a take-it-or-leave-it price set to be the higher of r_i and p_i^{VCG} .

Lemma 2.1. [18] *For any interdependent setting with matroid feasibility, the Generalized VCG auction with lazy reserve prices is ex post incentive compatible as long as for each i , the reserve price r_i is specified independently of s_i . For correlated, downward-closed settings, it is dominant-strategy incentive compatible.*

For completeness, we provide a brief proof sketch for this lemma.

PROOF SKETCH. It suffices to prove that GVCG without reserve prices is ex post IC, in other words, for all agents i , i is in the winning set W if and only if his value exceeds his threshold $p_i^{\text{VCG}}(\mathbf{s}_{-i})$. Recall that for a matroid constraint, the winning set W can be found by ordering agents by decreasing value and greedily selecting a maximal feasible set in this order. Also, in a matroid, the sizes of all maximal feasible subsets of a given set are equal. It follows then that whether or not an agent i belongs to the set W depends only on the set of agents preceding i in the ordering by decreasing value, and not on the relative ordering of those agents. Furthermore, if an agent in the winning set W unilaterally raises his signal, the single-crossing property implies that his rank in the ordering improves. Putting these two observations together we may conclude that the agent belongs to W as long as his signal exceeds his threshold s_i^* . \square

3 Correlated private values

For correlated private value single-item auctions, Ronen [28] proposed the lookahead auction which he showed 2-approximates the optimal revenue. In fact, the lookahead auction is simply VCG-L with conditional monopoly reserves, that is, where the reserve price for the highest bidder is set optimally based on the conditional distribution of the agent's value given others' values and the fact that the agent has the maximum value. We now give a natural extension of this result to matroid settings: we show that VCG-L with conditional monopoly reserves continues to give a 2-approximation to expected revenue in these settings.

Theorem 3.1. (Restatement of Theorem 1.2.) *The VCG-L auction with conditional monopoly reserves obtains at least half of the optimal revenue under a matroid feasibility constraint when agents have correlated private values.*

To prove Theorem 3.1, we will need the next well-known theorem on matroids. [See, e.g. 31, for a proof.]

Theorem 3.2. *Let B_1 and B_2 be any two independent sets of a matroid \mathcal{M} such that $|B_1| = |B_2|$. There exists a bijective mapping $g : B_1 \setminus B_2 \rightarrow B_2 \setminus B_1$ such that $\forall e \in B_1 \setminus B_2$, $B_2 \setminus \{e\} \cup \{g(e)\}$ is independent in \mathcal{M} .*

Proof of Theorem 3.1. We use VCG-L* to denote VCG-L with conditional monopoly reserves. Denote by Rev the expected revenue of the VCG-L* auction. The revenue of any optimal auction can be split into two parts, H and L : H is the expected revenue from W , the set of tentative winners, and L the revenue from the rest of the agents. Note that W here is random, determined by the realization of bidders' valuations, but H and L are expected values and not random. It suffices to show that Rev is no less than both H and L .

Rev is clearly at least H , since the VCG-L* auction runs the optimal auction for each agent in W , using all information available at that stage.

Let $W' \subseteq [n] \setminus W$ be an independent subset that maximizes the social welfare among bidders not in W . The expectation of $\sum_{j \in W'} v_j$ is an upper bound for L , since the auction cannot charge more than the agents' valuations. Therefore it suffices to show $\text{Rev} \geq \mathbf{E}[\sum_{j \in W'} v_j]$. Since $|W'| \leq |W|$, we can find a subset $U \subseteq W$ such that $U \cup W'$ is independent and $|U \cup W'| = |W|$. By Theorem 3.2 there exists a bijective mapping $g : W \setminus U \rightarrow W'$ such that for any bidder i in $W \setminus U$, $W \setminus \{i\} \cup \{g(i)\}$ is independent. Therefore, the VCG threshold p_i^{VCG} for each $i \in W \setminus U$ is at least $v_{g(i)}$. In the second stage of the VCG-L* auction, if the auctioneer simply sets VCG payment for each agent in W , a revenue of $\sum_{j \in W'} v_j$ would have been secured. By the optimality of the revenue from W in the lookahead auction, we have $\text{Rev} \geq \mathbf{E}[\sum_{j \in W'} v_j] \geq L$. \square

As we point out in Appendix A, it is not hard to see that, when the private valuations are independently drawn from regular distributions, the reserve prices optimally set for an agent in W will be either the VCG payment or the *monopoly reserve*, the optimal price one would set when selling a single item to this single bidder. In this case, Theorem 3.1 directly implies that the VCG-L auction with monopoly reserves 2-approximates the optimal auction, a result first shown by Dhangwatnotai et al. [12]. In Appendix A and Appendix B we also present similarly quick derivations of several other simple vs. optimal results from the literature.

The proof for Theorem 3.1, generalized from Ronen's proof for the lookahead auction, is direct and simple, without appealing to special properties of the distribution, nor any characterizations of the expected revenue (e.g. as virtual surplus). We use this as a convenient building block in the design and analysis of our mechanisms for interdependent value settings, which nonetheless call for new ideas to overcome the plain lookahead auctions' limitations.

4 General interdependent values

We now consider the general interdependent setting where agents' values depend on their own as well as on others' signals. We consider a natural generalization of the mechanism presented in the previous section for the correlated private values setting: generalized VCG-L with conditional monopoly reserves, or GVCG-L*. We saw an example in the introduction where the GVCG-L* mechanism does not always obtain a good approximation to the optimal revenue, even with just a single item for sale, simple value functions, and "nice" signal distributions. We get around the shortcomings in GVCG-L* by running the mechanism on a random subset of the agents. Importantly, we ensure that the agent with the highest value is left out of this subset with constant probability, so that the mechanism can extract revenue from agents with lower values as well. We

will now show that this modification is sufficient to obtain a constant factor approximation to optimal revenue.

Before we describe our mechanisms formally, we introduce some notation and prove a bound on the expected revenue of any ex post IC mechanism. Fix a vector of reported signals \mathbf{s} (which since we consider only incentive-compatible mechanisms are the true signals). We will use $v(A)$ to denote the total value of a subset A of agents: $\sum_{i \in A} v_i(\mathbf{s})$. Let W be the maximum value feasible set (e.g., in the case of a matroid constraint, an optimal base in the matroid). Let W' be the maximum value feasible set that is disjoint from W .

Lemma 4.1. *Let W and W' be defined as above. Then the expected revenue of any ex post IC mechanism is bounded by*

$$\mathbf{E} \left[\sum_{i \in W} R_i + v(W') \right]$$

where the expectation is over the randomness in signals; R_i is the expected revenue obtained by offering to serve agent i at the monopoly reserve price for the conditional value distribution of i , conditioned on \mathbf{s}_{-i} and on i being in W .

Proof. The revenue that any ex post IR mechanism can extract from agents not in W is at most $v(W')$. Therefore, we focus on the revenue that a mechanism can extract from the set W . This revenue is bounded from above by the sum of revenues of n ex post IC mechanisms, the i th one of which serves only agent i , and serves this agent only when this agent belongs to the set W . Any single-agent ex post IC mechanism is a (random) posted price mechanism with a posted price that depends on the reported signals of the other agents; therefore, the revenue obtained by such a mechanism is at most R_i . \square

4.1 Single item setting

We now describe our mechanism for the single item setting. Consider the following mechanism that we call randomized GVCG-L*:

1. Ask each agent to report his signal.
2. Include each agent in a set Z with probability $2/3$.
3. Run GVCG-L* on Z ; when determining the reserve price for agent i , condition on signals of all agents except i (including those not in Z).

Theorem 4.2. *Consider a single item setting with interdependent values, where the value functions satisfy assumptions (A.1)–(A.3). The randomized GVCG-L* is ex post IC and achieves a 4.5-approximation to the optimal ex post IC mechanism.*

Proof. Let $\mathbf{s} = (s_1, s_2, \dots, s_n)$ denote the agents' signals. Without loss of generality, suppose that $v_1(\mathbf{s}) \geq v_2(\mathbf{s}) \geq v_j(\mathbf{s})$ for all $j > 2$. Define

$$s_1^* = \operatorname{argmin}_s \{v_1(s, \mathbf{s}_{-1}) \geq v_j(s, \mathbf{s}_{-1}) \quad \forall j > 1\},$$

In other words, s_1^* is the smallest signal at which agent 1's value is the highest, given \mathbf{s}_{-1} . Let agent i be the one who determines the threshold s_1^* , that is, $v_1(s_1^*, \mathbf{s}_{-1}) = v_i(s_1^*, \mathbf{s}_{-1})$. Note that i may or may not be the agent with the second highest value at \mathbf{s} .

Lemma 4.1 implies that

$$OPT \leq R_1 + v_2(s_1, s_{-1})$$

where R_1 is the expected revenue from agent 1 conditioned on $s_1 \geq s_1^*$, and \mathbf{s}_{-1} .

On the other hand, the expected revenue of randomized GVCGL* is at least

$$Pr(1 \in Z, i \in Z) \cdot R_1 + Pr(2 \in Z, 1 \notin Z) \cdot R_2 = \frac{2}{9}(2R_1 + R_2)$$

where R_2 is the expected revenue from agent 2 conditioned on the fact that $v_2(\mathbf{s})$ is highest among agents in Z , and conditioned on \mathbf{s}_{-2} . Therefore, randomized GVCGL* recovers within a small factor the first component of the optimal revenue.

We will now bound the second term in OPT , namely $v_2(s_1, \mathbf{s}_{-1})$, in terms of R_1 and R_2 . Observe that fixing \mathbf{s}_{-1} , agent 1's value at his threshold s_1^* is at least as large as any agent's value. Furthermore, R_1 extracts at least this threshold value. That is,

$$v_2(s_1^*, \mathbf{s}_{-1}) \leq v_1(s_1^*, \mathbf{s}_{-1}) \leq R_1.$$

Next, using assumption (A.3) we get that

$$v_2(s_1, \mathbf{s}_{-1}) - v_2(s_1^*, \mathbf{s}_{-1}) \leq v_2(s_1, 0, \mathbf{s}_{-12}) - v_2(s_1^*, 0, \mathbf{s}_{-12}) \leq v_2(s_1, 0, \mathbf{s}_{-12}) \leq R_2$$

where the third inequality follows from noting that agent 2's value conditioned on others' signals is at least $v_2(s_1, 0, \mathbf{s}_{-12})$. Therefore, we have $v_2(s_1, \mathbf{s}_{-1}) \leq R_1 + R_2$, and $OPT \leq 2R_1 + R_2$. \square

4.2 Matroid setting

Next we consider settings with a matroid feasibility constraint.⁶ We will run the GVCGL* mechanism on a random subset Z of agents. GVCGL* will select the maximum value feasible subset of this set, call it T , and optimally price the agents in that set conditioned on others' signals. Define $T' = T \setminus W$. Our goal in selecting Z is two-fold. First, we want to ensure that T contains agents in W with high probability, so that we can recover (to some approximation) the R_i 's for agents i in W . Second, we want to ensure that $T' = T \setminus W$ contains value comparable to that in W' . We will then be able to “charge” the value of every agent in T' to the revenue we recover from this agent plus the revenue we recover from some agent in W . With these goals in mind, we define the mechanism as follows.

1. Ask each agent to report his signal.
2. With probability $1/2$ let Z be the set of all agents; and with probability $1/2$ include each agent independently in Z with probability $1/2$.
3. Run GVCGL* on Z ; when determining the reserve price for an agent i , condition on the signals of all agents except i (including those not in Z).

⁶We note that our approach also applies to knapsack constraints with the modification that in GVCGL* instead of picking the maximum value feasible set in the first step we greedily pick a maximal feasible set.

The following two lemmas capture the properties that we require from T and T' .

Lemma 4.3. $\mathbf{E}[v(T')] \geq \frac{1}{8} \mathbf{E}[v(W')]$.

Lemma 4.4. *For any random T' there exists an injective mapping $f_{T'}$ from the elements of T' to the elements of W such that for each $j \in T'$ and $i = f_{T'}(j)$, we have,*

$$v_j(s_i^*, \mathbf{s}_{-i}) \leq v_i(s_i^*, \mathbf{s}_{-i}).$$

Before proving the lemmas, we give a proof of our main theorem:

Theorem 4.5. *Consider an interdependent values setting with a matroid feasibility constraint where the value functions satisfy assumptions (A.1)–(A.3). The randomized GVCG-L* is ex post IC and achieves a 18-approximation to the optimal ex post IC mechanism.*

Proof. Recall that for an agent $i \in W$, R_i denotes the expected revenue we can obtain from this agent by setting an optimal reserve price for him conditioned on the other agents' signals and on being in the set W . With probability 1/2 our mechanism runs GVCG-L* on all of the agents, and therefore obtains this revenue R_i from every agent $i \in W$. With the remaining probability, the mechanism obtains some revenue from agents in W and some revenue from agents in $T' = T \setminus W$. For agents j not in W , we can lower bound the revenue that the mechanism can collect from them if they end up in T' : this is the expected revenue obtained by optimally pricing agent j conditioned on \mathbf{s}_{-j} ; ⁷ call this quantity \tilde{R}_j .

We therefore get that the expected revenue of the mechanism is at least

$$\frac{1}{2} \sum_{i \in W} R_i + \mathbf{E} \left[\sum_{j \in T'} \tilde{R}_j \right].$$

On the other hand, recall that by Lemma 4.1 we can bound the optimal revenue as

$$OPT \leq \sum_{i \in W} R_i + v(W').$$

We will now upperbound the second term in terms of the R_i 's and \tilde{R}_j 's. Fix a random choice of T' . Let $j \in T'$ and $i = f_{T'}(j)$, as defined in Lemma 4.4. Then

$$\tilde{R}_j \geq v_j(0, \mathbf{s}_{-j}) \geq v_j(s_i, 0, \mathbf{s}_{-ij}) - v_j(s_i^*, 0, \mathbf{s}_{-ij}) \geq v_j(s_i, \mathbf{s}_{-i}) - v_j(s_i^*, \mathbf{s}_{-i})$$

where the last inequality follows from assumption (A.3). But by Lemma 4.4, for that same $i = f_{T'}(j)$,

$$R_i \geq v_i(s_i^*, \mathbf{s}_{-i}) \geq v_j(s_i^*, \mathbf{s}_{-i}),$$

so

$$\tilde{R}_j + R_i \geq v_j(\mathbf{s}).$$

⁷Note that in this lower bound we do not condition on j being in set T' ; the extra conditioning can potentially increase the expected revenue, but this increase is difficult to estimate because it depends on the random set Z .

Finally, by Lemma 4.3,

$$\begin{aligned} v(W') &\leq 8 \mathbf{E} \left[\sum_{j \in T'} v_j(\mathbf{s}) \right] \\ &\leq 8 \mathbf{E} \left[\sum_{j \in T'} (\tilde{R}_j + R_{f_{T'}(j)}) \right]. \end{aligned}$$

Putting it together:

$$\begin{aligned} OPT &\leq \sum_{i \in W} R_i + v(W') \\ &\leq \sum_{i \in W} R_i + 8 \mathbf{E} \left[\sum_{j \in T'} (\tilde{R}_j + R_{f_{T'}(j)}) \right] \\ &\leq 9 \sum_{i \in W} R_i + 8 \mathbf{E} \left[\sum_{j \in T'} \tilde{R}_j \right], \end{aligned}$$

and we get an 18-approximation. \square

It remains to prove the lemmas, for which we will need the following fact whose proof can be found e.g., in Schrijver [31].

Lemma 4.6 (Strong Basis Exchange). *Let B and B' be two bases of a matroid. Then for all $x \in B \setminus B'$, there is a $y \in B' \setminus B$ such that both $B - x + y$ and $B' - y + x$ are bases.*

PROOF OF LEMMA 4.3. In step (2) with probability $1/2$ we construct Z by sampling each agent. Let us condition on this event. We will fix a particular order in which to sample elements for Z and in the process inductively update a particular independent set. It will end up containing all the elements of $W \cap Z$, and each element of W' will be in it, and in Z , with probability $1/4$.

Assume that $|W'| = |W|$. (If not, extend W' to a full base using elements of W and adapt the following arguments appropriately.) Fix an ordering on the elements of W , say a_1, \dots, a_r . We will also be considering the elements of W' in a particular order b_1, \dots, b_r to be determined.

As we sample elements for Z , we will inductively update two (ordered) sets of matroid elements A_i and B_i that satisfy the following invariants:

- A_i and B_i are both bases.
- The sampling has been done only for elements $\{a_1, \dots, a_i\}$ and $\{b_1, \dots, b_i\}$.
- The first i elements in A_i and B_i are the same.
- All elements in $\{a_1, \dots, a_i\} \cap Z$ are in A_i .
- Each element in $\{b_1, \dots, b_i\}$ is in $Z \cap A_i$ with probability $1/4$.

Initially $A_0 = W$ and $B_0 = W'$ satisfy the above invariants. Now suppose that A_{i-1} and B_{i-1} have been constructed and satisfy the above invariants. Apply the strong basis exchange lemma to find an element $b_i \in B_{i-1}$ such that $A_{i-1} - a_i + b_i$ and $B_{i-1} - b_i + a_i$ are both bases.

Now toss the coin to determine if $a_i \in Z$. If so, $A_i := A_{i-1}$ and $B_i := B_{i-1} - b_i + a_i$. If $a_i \notin Z$, then $A_i := A_{i-1} - a_i + b_i$ and $B_i := B_{i-1}$. Finally toss the coin to determine if $b_i \in Z$. It is easy to see that all the required invariants are satisfied.

At the end of the process, consider the set $A^* = A_r \cap Z$ (where r is the rank of the matroid). Note that all of the agents in $W \cap Z$ belong to A_r , so they also belong to A^* . On the other hand, the greedy algorithm for matroids includes all agents in $W \cap Z$ into T . So, $A^* \setminus W$ is a candidate for the set $T' = T \setminus W$. We now note that $A^* \setminus W$ contains sufficient value: by the invariant above, every element of W' is in $A_r \cap Z$ with probability $1/4$. Therefore we get:

$$\mathbf{E} [v(T')] \geq \mathbf{E} [v(W' \cap A^*)] = \frac{1}{4}v(W').$$

Removing the conditioning specified earlier yields the lemma. \square

PROOF OF LEMMA 4.4. Let the rank of the matroid be r . Recall that $|T| \leq r$ and $T' = T \setminus W$. For now assume that T' is full rank. We construct the function $f_{T'}$ by showing that there is a matching between W and T' such that for each i in W , its match $j = M(i) \in T'$ has $v_j(s_i^*, \mathbf{s}_{-i}) \leq v_i(s_i^*, \mathbf{s}_{-i})$.

To this end, construct a bipartite graph $G_{W,T'}$ where there is an edge from $i \in W$ to $j \in T'$ if $v_j(s_i^*, \mathbf{s}_{-i}) \leq v_i(s_i^*, \mathbf{s}_{-i})$.

Let G_i denote the neighbors of $i \in W$ in the graph, and $B_i = T' \setminus G_i$. We will use Hall's theorem to show that there is a perfect matching in the graph. Specifically, for any $S \subseteq W$, we will show that $|\Gamma(S)| \geq |S|$ where $\Gamma(S)$ are the neighbors of S .

Without loss of generality assume that the agents of W are indexed in order of decreasing value at \mathbf{s} (so their indices are $1, \dots, r$). For all $k \leq r$, and $j < k$, we have $v_j(s_k^*, \mathbf{s}_{-k}) \geq v_k(s_k^*, \mathbf{s}_{-k})$. This is because $v_j(\mathbf{s}) \geq v_k(\mathbf{s})$, and the single-crossing condition ensures that agent k cannot be above j at s_k^* and end up at or below j at s_k^* .

For each i , let Y_i be the set of agents in $W \cup T'$ which have greater value than i at (s_i^*, \mathbf{s}_{-i}) . Since s_i^* is critical for i ,

$$\text{rank}(Y_i) < \text{rank}(Y_i \cup \{i\}). \quad (1)$$

Now, let $S = \{i_1, \dots, i_k\}$ (ordered in decreasing order of value at \mathbf{s}), let $B = \cap_{i \in S} B_i$, and let S_j be the first $j-1$ elements of S . Then for each element i_j of S , by submodularity of the rank function, the fact that $S_j \cup B \subseteq Y_{i_j}$ and (1), we have

$$\text{rank}(S_j \cup B) < \text{rank}(S_j \cup B \cup \{i_j\}) \quad \text{and thus} \quad \text{rank}(S_j \cup B \cup \{i_j\}) = \text{rank}(S_j \cup B) + 1.$$

Induction on j then implies that $\text{rank}(S \cup B) = |S| + |B|$. Since this is at most r , $|B| \leq r - |S|$, and so $|\Gamma(S)| = |T'| - |B| = r - |B| \geq |S|$.

To handle the case that T' is not full rank, simply augment it to a full basis using elements of W , and show that there is a matching between $W \setminus T'$ and T' by following the same argument. \square

5 Conclusions and discussion

In the previous sections we showed that GVCG-L* provides constant factor approximations to the optimal ex post IC revenue in interdependent settings under certain assumptions. There are several directions for further work which we now discuss.

- *Prior independence.* In independent values settings it has been shown that under certain assumptions (e.g., regularity of the value distributions) it is possible to design mechanisms that do not require knowing the value distributions and yet obtain an approximation to the optimal expected revenue. Such mechanisms are called prior-independent. One way of designing a prior-independent mechanism, called the single-sample approach, is to use a reserve price based mechanism (such as GVCG-L) and replace the reserve price by an independent sample from the agent’s value distribution. The mechanisms that we design require knowing the conditional distributions of agents’ values in order to determine an appropriate reserve price. Is it possible to design prior-independent mechanisms in interdependent settings?

It is easy to see that assuming regularity of the conditional distribution a single-sample approach works: we can replace the optimal reserve price for an agent in GVCG-L by a random independent draw from the agent’s value distribution conditioned on others’ reported signals, and still obtain a constant factor approximation to the optimal expected revenue. However, this approach is unsatisfying. In independent value settings, if there are several agents with identically distributed values, we can remove one of these agents at random from the auction and use his value as the random reserve. In interdependent value settings, we cannot truly “remove” an agent from the auction because her signal affects others’ values⁸.

Another approach to prior-independent mechanism design is to artificially limit supply [11, 30]. In the interdependent values context, we might ask, for instance: when agents are symmetric and conditional value distributions satisfy regularity, does the VCG mechanism with $k/2$ units for sale approximate the expected revenue of an optimal k -unit auction for $k \geq 2$?

- *Relaxing the assumptions.* In the absence of the single-crossing assumption, the VCG mechanism and its variants are no longer ex post incentive compatible. While it appears to be challenging to characterize the optimal mechanism in this context, approximations may be tractable. Likewise, beyond matroid feasibility constraints, VCG and its variants are no longer necessarily ex post IC. Is it possible to approximate the optimal mechanism in this case?
- *Ex post IC versus Bayesian IC.* In independent value settings with single-parameter agents, Myerson shows that ex post IC mechanisms are as strong as BIC mechanisms in terms of extracting revenue. In interdependent settings, Cremer and McLean show that BIC mechanisms can extract the entire social welfare under mild assumptions, which can in general be much larger than the optimal ex post IC revenue. However, Cremer and McLean’s mechanism violates ex post individual rationality. Can BIC mechanisms obtain more revenue than ex post IC mechanisms under an ex post IR constraint? Can we approximate this revenue?

⁸Roughgarden and Talgam-Cohen [29] show that under the so-called Lopomo assumptions, which includes strong MHR and symmetry across agents, one can get a constant factor approximation by “reusing” the signal of an agent to determine a random reserve without dropping the agent from the auction.

References

- [1] Abraham, I., Athey, S., Babaioff, M., and Grubb, M. (2013). Peaches, lemons, and cookies: designing auction markets with dispersed information. In *ACM Conference on Electronic Commerce*.
- [2] Ausubel, L. (1999). A generalized vickrey auction. *Econometrica*.
- [3] Azar, P., Daskalakis, C., Micali, S., and Weinberg, S. M. (2013). Optimal and efficient parametric auctions. In *24th ACM-SIAM Symposium on Discrete Algorithms*.
- [4] Babaioff, M., Kleinberg, R., and Leme, R. P. (2012). Optimal mechanisms for selling information. In *ACM Conference on Electronic Commerce*.
- [5] Chung, K. and Ely, J. (2006). Ex-post incentive compatible mechanism design. Unpublished manuscript.
- [6] Chung, K.-S. and Ely, J. C. (2007). Foundations of dominant-strategy mechanisms. *The Review of Economic Studies*, 74(2):447–476.
- [7] Clarke, E. H. (1971). Multipart pricing of public goods. *Public Choice*, 2:19–33.
- [8] Crémer, J. and McLean, R. (1985). Optimal selling strategies under uncertainty for a discriminating monopolist when demands are interdependent. *Econometrica*, 53(2).
- [9] Crémer, J. and McLean, R. (1988). Full extraction of the surplus in bayesian and dominant strategy auctions. *Econometrica*, 56(6).
- [10] Dasgupta, P. and Maskin, E. (2000). Efficient auctions. *The Quarterly Journal of Economics*, 115(2):341–388.
- [11] Devanur, N., Hartline, J., Karlin, A., and Nguyen, T. (2011). Prior-independent multi-parameter mechanism design. In *Internet and Network Economics*, pages 122–133. Springer.
- [12] Dhangwatnotai, P., Roughgarden, T., and Yan, Q. (2010). Revenue maximization with a single sample. In *ACM Conference on Electronic Commerce*, pages 129–138.
- [13] Dobzinski, S., Fu, H., and Kleinberg, R. D. (2011). Optimal auctions with correlated bidders are easy. In *STOC*, pages 129–138.
- [14] Groves, T. (1973). Incentives in Teams. *Econometrica*, 41:617–631.
- [15] Hartline, J. D. and Roughgarden, T. (2009). Simple versus optimal mechanisms. In *ACM Conference on Electronic Commerce*, pages 225–234.
- [16] Klemperer, P. (1998). Auctions with almost common values: The ‘wallet game’ and its applications. *European Economic Review*, 42(3-5):757–769.
- [17] Krishna, V. (2010). *Auction Theory*. Academic Press.
- [18] Li, Y. (2013). Approximation in mechanism design with interdependent values. In *ACM Conference on Electronic Commerce*, pages 675–676.

- [19] Lopomo, G. (2000). Optimality and robustness of the english auction. *Games and Economic Behavior*, 36:219–240.
- [20] McAfee, R. P., McMillan, J., and Reny, P. J. (1989). Extracting the surplus in the common-value auction. *Econometrica*, pages 1451–1459.
- [21] McAfee, R. P. and Reny, P. J. (1992). Correlated information and mechanism design. *Econometrica*, 60(2):395–421.
- [22] Milgrom, P. (2004). *Putting Auction Theory to Work*. Cambridge University Press.
- [23] Milgrom, P. and Weber, R. J. (1982). A theory of auctions and competitive bidding. *Econometrica*, 50:1089–1122.
- [24] Myerson, R. (1981). Optimal auction design. *Mathematics of Operations Research*, 6(1):pp. 58–73.
- [25] Neeman, Z. (2003). The effectiveness of english auctions. *Games and Economic Behavior*, 43(2):214–238.
- [26] Papadimitiou, C. and Pierrakos, G. (2011). On optimal single-item auctions. In *STOC*, pages 119–128.
- [27] Rahman, D. (2014). Surplus extraction on arbitrary type spaces. Forthcoming in *Theoretical Economics*.
- [28] Ronen, A. (2001). On approximating optimal auctions. In *ACM Conference on Electronic Commerce*, pages 11–17.
- [29] Roughgarden, T. and Talgam-Cohen, I. (2013). Optimal and near-optimal mechanism design with interdependent values. In *ACM Conference on Electronic Commerce*, pages 767–784.
- [30] Roughgarden, T., Talgam-Cohen, I., and Yan, Q. (2012). Supply-limiting mechanisms. In *ACM Conference on Electronic Commerce*, pages 844–861.
- [31] Schrijver, A. (2003). *Combinatorial Optimization : Polyhedra and Efficiency (Algorithms and Combinatorics)*. Springer.
- [32] Segal, I. (2003). Optimal pricing mechanisms with unknown demand. *American Economic Review*, 93.
- [33] Syrgkanis, V., Kempe, D., and Tardos, E. (2013). Information asymmetries in common-value auctions with discrete signals. In *SSRN*.
- [34] Vickrey, W. (1961). Counterspeculation, Auctions and Competitive Sealed Tenders. *Journal of Finance*, pages 8–37.
- [35] Vohra, R. (2011). *Mechanism Design: A Linear Programming Approach*. Cambridge University Press.
- [36] Wilson, R. (1969). Competitive bidding with disparate information. *Management Sciences*, 15.

A Simple vs. Optimal Results via Lookahead Auctions

In this appendix we give quick derivations for several results in the literature showing that mechanisms simple in form approximates the optimal revenue.

Recall that a *monopoly reserve price* for a bidder (or rather, her value distribution) is the optimal take-it-or-leave-it price that an auctioneer should set when selling the item to the bidder alone. For bidder i , we denote her monopoly reserve price as r_i^{mon} .

Theorem A.1. *In general single-dimensional auction settings, when bidders' private valuations are drawn independently from regular distributions, the VCG-L auction with conditional monopoly reserve prices is identical to the VCG-L auction with (unconditional) monopoly reserve prices.*

Proof. In the lookahead auction, once a bidder i is in W , the auctioneer runs an optimal auction for the bidder conditioning on the other bidders' values and the fact $i \in W$ given these values. In the independent value settings, this simply amounts to setting an optimal take-it-or-leave-it price for the bidder conditioning on $v_i \geq p_i^{\text{VCG}}$. We show that this optimal price is equal to $\max\{p_i^{\text{VCG}}, r_i^{\text{mon}}\}$. With this, a quick comparison would confirm that the lookahead auction is identical with VCG-L with monopoly reserves.

The distribution of v_i , conditioning on $v_i \geq p_i^{\text{VCG}}$ is just F_i truncated at p_i^{VCG} , i.e., $F_i^{\text{cond}}(v) = \frac{F_i(v) - F_i(p_i^{\text{VCG}})}{1 - F_i(p_i^{\text{VCG}})}$, for any $v \geq p_i^{\text{VCG}}$, and $F_i^{\text{cond}}(v) = 0$ for any $v < p_i^{\text{VCG}}$. Setting any price below p_i^{VCG} would be obviously suboptimal, and the revenue of setting a price of $p \geq p_i^{\text{VCG}}$ is

$$p(1 - F_i^{\text{cond}}(p)) = p \cdot \left(1 - \frac{F_i(p) - F_i(p_i^{\text{VCG}})}{1 - F_i(p_i^{\text{VCG}})}\right) = p(1 - F_i(p)) \cdot \frac{1}{1 - F_i(p_i^{\text{VCG}})}.$$

In words, for prices above p_i^{VCG} , the conditional expected revenue is simply scaled up by a constant factor of $\frac{1}{1 - F_i(p_i^{\text{VCG}})}$ from the unconditioned distribution. Therefore, if $p_i^{\text{VCG}} < r_i^{\text{mon}}$, the optimal price to set for the conditional distribution is still r_i^{mon} ; and if $p_i^{\text{VCG}} \geq r_i^{\text{mon}}$, by regularity, the expected revenue monotonically decreases as the price rises above p_i^{VCG} ; therefore the optimal price to set is p_i^{VCG} itself. \square

Given Theorem 3.1, this immediately implies the following result by Dhangwatnotai et al. [12].

Corollary A.2 (12). *In matroid settings where bidders' valuations are drawn independently from regular distributions, the VCG-L auction with monopoly reserve prices gives at least half of the optimal revenue.*

As another illustration of the usefulness between the lookahead auction and the reserve-based VCG auctions, we use the analysis of the lookahead auction to give a short re-derivation for a result by Azar et al. [3], a major building block in that work.

Let $R_i(p)$ denote the expected revenue from bidder i by setting a price of p .

Theorem A.3 (Theorem 3.1 in 3). *For each bidder i , let r_i be a price drawn randomly from a certain distribution, such that in the single-bidder setting, $\mathbf{E}_{r_i}[R_i(r_i)] \geq \alpha R_i(r_i^{\text{mon}})$. Then in all downward closed settings where bidders' valuations are drawn independently from regular distributions, VCG-L with random reserve prices (r_1, \dots, r_n) gets at least α -fraction of the revenue of VCG-L with monopoly reserve prices (VCG-LM).*

Proof. Fixing a tentative winner $i \in W$, and condition on the VCG price p_i^{VCG} . For a reserve price r_i , let ρ denote the ratio $R_i(r_i)/R_i(r_i^{\text{mon}})$, then we have $\mathbf{E}[\rho] \geq \alpha$. Let ρ^{cond} denote the ratio between the expected VCG-L revenue from using reserve price r_i and that from VCG-LM, conditioning on $v_i \geq p_i^{\text{VCG}}$. It suffices to show $\mathbf{E}[\rho^{\text{cond}}] \geq \alpha$. We will show $\rho^{\text{cond}} \geq \rho$ pointwise. Since, as we have shown in the proof of ??, the revenue from all prices above p_i^{VCG} is scaled up by a factor of $1/(1 - F_i(p_i^{\text{VCG}}))$ when conditioning on $v_i > p_i^{\text{VCG}}$, for $r_i \geq p_i^{\text{VCG}}$ we have

$$\rho^{\text{cond}} = \frac{R_i(r_i)/(1 - F_i(p_i^{\text{VCG}}))}{R_i(\max\{p_i^{\text{VCG}}, r_i^{\text{mon}}\})/(1 - F_i(p_i^{\text{VCG}}))} \geq \frac{R_i(r_i)}{R_i(r_i^{\text{mon}})} = \rho.$$

For $r_i < p_i^{\text{VCG}}$, if $p_i^{\text{VCG}} \geq r_i^{\text{mon}}$, then both auctions will use p_i^{VCG} , and $\rho^{\text{cond}} = 1 \geq \rho$; if $p_i^{\text{VCG}} < r_i^{\text{mon}}$, we have

$$\rho^{\text{cond}} = \frac{R_i(p_i^{\text{VCG}})/(1 - F_i(p_i^{\text{VCG}}))}{R_i(r_i^{\text{mon}})/(1 - F_i(p_i^{\text{VCG}}))} = \frac{R_i(p_i^{\text{VCG}})}{R_i(r_i^{\text{mon}})} \geq \frac{R_i(r_i)}{R_i(r_i^{\text{mon}})} = \rho.$$

The last inequality comes from regularity, i.e., $r_i < p_i^{\text{VCG}} \leq r_i^{\text{mon}}$ implies $R_i(r_i) \leq R_i(p_i^{\text{VCG}})$. This completes the proof. \square

B VCG-L vs. VCG-E

In this appendix we compare the revenue of the VCG-E and VCG-L auctions with the same set of reserves in for independent settings with regular distributions. Although we end the appendix with Corollary B.3, a known result from [?], the general comparison we make here concerns VCG-L and VCG-E auctions with reserves that are not necessarily the monopoly reserves; in fact, even for the monopoly reserve prices, the revenue dominance of VCG-E that we show here is not known in the literature. Given the practicality and popularity of such reserve-based auctions, we believe this comparison to be of independent interest.

Recall that the VCG-E is the auction that first removes all bidders bidding below their reserve prices and then run VCG-L on the remaining bidders. It is not too hard to see that adding the first stage improves the social welfare, but the revenue analysis is nontrivial. We need to analyze two counteractive aspects. On the one hand, VCG-E removes bidders that are not going to afford their take-it-or-leave-it price later in the auction, and helps free up resources to be sold to the other bidders; on the other hand, the existence of more bidders always creates more competition, and makes the VCG payment for other bidders higher and therefore helps create revenue. Theorem B.2 shows that, when the reserve prices are monopoly reserves, the first force is the dominant one.

The analysis is not trivial, and we first present the case of single-item auctions as a warm-up. Throughout this appendix we denote by r_i^{mon} the monopoly reserve price for a bidder i , i.e., the optimal posted price one would set when selling a single item to this bidder.

Theorem B.1. *In a single item auction where bidders' valuations are drawn independently from regular distributions, the second price auction with eager monopoly reserve prices (VCG-EM) generates weakly more revenue than the second price auction with lazy monopoly reserve prices (VCG-LM).*

Proof. The key idea of the proof is to consider the expected revenue from each bidder, conditioning on the other bidders' valuations. We calculate the conditional expected revenues from the two auctions and compare them. The computation of the conditional revenue does not assume conditions

on the reserve price or the valuation distribution; only in the last step do we use the property of monopoly reserves and the valuation distribution's regularity.

Fix a bidder i , we condition on all other bidders' valuations. Denote by h_{-i}^L the highest valuation among these other bidders, i.e., $h_{-i}^L = \max_j \{v_j \mid j \neq i\}$. Denote by h_{-i}^E the highest valuation among bidders except i who bid above their reserve prices, i.e., $h_{-i}^E = \max_j \{v_j \mid j \neq i, v_j \geq r_j\}$. Clearly, $h_{-i}^L \geq h_{-i}^E$.

When $h_{-i}^L \leq r_i$, the revenue extracted from bidder i is exactly the same in VCG-E and VCG-L, because in both auctions, bidder i will make a payment if and only if his bid is at least r_i , in which case he pays r_i . Therefore we need only consider the case $h_{-i}^L > r_i$.

In VCG-L, bidder i makes a payment of h_{-i}^L if and only if his valuation is above h_{-i}^L (since $h_{-i}^L > r_i$). The expected revenue from him is then $h_{-i}^L(1 - F_i(h_{-i}^L))$. In VCG-E, bidder i makes a payment if and only if his valuation is above both r_i and h_{-i}^E , and the payment is $p_i^E = \max\{r_i, h_{-i}^E\}$. The expected revenue from i is therefore $p_i^E(1 - F_i(p_i^E))$. Note that $h_{-i}^L \geq \max\{r_i, h_{-i}^E\} = p_i^E$.

But the revenues $h_{-i}^L(1 - F_i(h_{-i}^L))$ and $p_i^E(1 - F_i(p_i^E))$ are simply the expected revenue we get by setting a price of h_{-i}^L or p_i^E to bidder i , respectively. By regularity of the valuation distribution, the expected revenue monotonically decreases as we raise the price above the monopoly reserve. Since $h_{-i}^L \geq p_i^E \geq r_i = r_i^{\text{mon}}$, we have

$$h_{-i}^L(1 - F_i(h_{-i}^L)) \leq p_i^E(1 - F_i(p_i^E)).$$

The above is a conditional analysis, but we see that the expected revenue in VCG-LM from each bidder i is no more than the expected revenue in VCG-EM from the same bidder, no matter what the other bidders bid. Our theorem immediately follows. \square

Theorem B.2. *For independent private value settings with matroid feasibility constraints and regular valuation distributions, the VCG-E auction has weakly better social welfare than VCG-L, and when the reserves are at least the monopoly reserves, the VCG-E auction also obtains weakly more revenue than VCG-L.*

Proof. As in the proof of Theorem B.1, we focus on the expected revenue from a fixed bidder i , conditioning on all other bidders' valuations. Let $p_i^{\text{VCG-L}}$ and $p_i^{\text{VCG-E}}$ denote the VCG payment for bidder i in VCG-L and VCG-E, respectively. (Note that the actual payment to be made by bidder i in these auctions is the higher of the reserve r_i and the VCG threshold; but we will focus mostly on the VCG thresholds in this proof.) The key step is to show $p_i^{\text{VCG-L}} \geq p_i^{\text{VCG-E}}$. Recall that in VCG-E, bidders with valuations below their reserve prices are excluded from both the allocation and payment calculation. Effectively, such bidders' bids are replaced by 0 in the calculation of VCG payments. This constitutes the only difference in computing $p_i^{\text{VCG-L}}$ and $p_i^{\text{VCG-E}}$. It therefore suffices to show that bidder i 's VCG payment weakly decreases as other bidders' bids are zeroed out.

We show this by expressing $p_i^{\text{VCG-L}}$ and $p_i^{\text{VCG-E}}$ in terms of a submodular function. Define a set function $f : 2^{[n]} \rightarrow \mathbb{R}$ by

$$f(S) = \max_{T \subseteq S, T \in \mathcal{I}} \sum_{j \in T} v'_j,$$

where $v'_j = v_j$ for $j \neq i$ and $v'_i = \sum_{j \neq i} v_j + 1$. Then $p_i^{\text{VCG-L}} = f([n] \setminus \{i\}) - (f([n] - v'_i))$. This is because $f([n] \setminus \{i\})$ represents the social welfare of the other bidders if i were not present, and

$f([n]) - v'_i$ represents the social welfare of the other bidders if i were to be a winner (v'_i is set high enough such that the optimal set has to include i). The difference between the two is then the externality that i imposes on the other bidders by his winning, and therefore is equal to the VCG threshold p_i^{VCGL} . Similarly, let U denote the set of bidders whose valuations are above their reserve prices, i.e., $U = \{j \mid v_j \geq r_j\}$, then $p_i^{\text{VCGE}} = f(U) - (f(U \cup \{i\}) - v'_i)$.

f is defined to be the result of a linear maximization over a matroid, and is well known to be submodular. Therefore, as $U \subseteq [n]$, $f([n]) - f([n] \setminus \{i\}) \leq f(U \cup \{i\}) - f(U)$. This gives $p_i^{\text{VCGL}} \geq p_i^{\text{VCGE}}$.

With this, the proof to Theorem B.1 easily generalizes by replacing h_{-i}^{L} and h_{-i}^{E} with p_i^{VCGL} and p_i^{VCGE} , respectively, and we omit the rest of the proof. \square

Corollary B.3 (15). *In matroid settings where bidders' valuations are drawn independently from regular distributions, the VCG-E auction with monopoly reserve prices gives at least half of the optimal revenue.*

Remark 1. For more general environments (even downward closed), it need not be true that p_i^{VCGL} is at least p_i^{VCGE} . For example, consider $\mathcal{I} = \{\emptyset, \{1\}, \{2\}, \{1, 2\}, \{3\}\}$. In this feasibility system, bidder 1's VCG payment weakly increases when v_2 decreases. So p_1^{VCGL} can be lower than p_1^{VCGE} .